In the Specification:

Please amend the paragraph beginning on Page 2, line 13 and ending on Page 3, line 11, as follows:

The number of objects involved in servicing a content request may range from a single stored object to a relatively large number of objects (often, on the order of tens of objects). (The terms "stored object" and "object" are used interchangeably herein to refer to an object or file which is stored on a storage medium--or which may, in some cases, be distributed across more than one storage medium. It should be noted that references herein to objects are not to be construed as limiting the present invention to the field of object-oriented programming. Furthermore, the terms "content" and "document content" as used herein are intended to be synonymous with one or more objects or files unless the reference context indicates otherwise.) Studies show that requests for objects generally follow a very distinct statistical distribution pattern. That is, in an average content distribution environment, there are typically a large number of requests for a relatively small number of objects, while there are a smaller number of requests for a much larger population of the objects. This distribution pattern for content requests generally follows what is known as a "Zipf distribution" (or a "Zipf-like" distribution). A Zipf distribution is a particular type of distribution pattern, wherein a double-logarithmic plotting of the distribution follows a straight line. Papers discussing the Zipf-like distribution of content requests include "Zipf Curves and Website Popularity", published on the Internet at location ~~http://www.useit.com/alertbox/zipf.html~~ useit.com/alertbox/zipf.html (Apr. 15, 1997) and "Do Websites Have Increasing Returns", Jakob Nielsen, published on the Internet at location ~~http://www.useit.com/alertbox/9704b.html~~ useit.com/alertbox/9704b.html (Apr. 15, 1997).

Please amend the paragraph beginning on Page 3, line 19 and ending on Page 4, line 16, as follows:

FIG. 1 provides a diagram of a representative server site 100 in which a content request is serviced. (The term "server site" as used herein refers to a collection of server nodes that serve Web content associated with a given fully-qualified domain name. For example, the server site 100 in FIG. 1 may, for purposes of example, serve content for a domain name such as ~~"www.ibm.com"~~ ibm.com.) In this example, a content request 110 is transmitted from a client (not shown) through a network such as the Internet 120 and then to

a load balancing host 130 (that is, a computing device which distributes incoming requests across a plurality of Web servers 140 to balance the processing load). The load balancing host 130 may then select one of the Web servers 140 (such as Apache, Netscape®, or Microsoft® servers), according to the load balancing strategy which has been implemented in host 130. To serve the requested content, a particular Web server may invoke the services of an application server (such as a WebSphere application server which is available from the International Business Machines Corporation, or "IBM"), where this application server may be co-located with the Web server 140 in a single hardware box or may be located at a different device 150. The Web server may also or alternatively invoke the services of a back-end enterprise data server 160 (such as an IBM OS/390® server running the DB/2, CICS®, and/or MQI products from IBM), which may in turn access one or more databases 170 or other data repositories. ("WebSphere", "OS/390", and "CICS" are registered trademarks of IBM.)

Please amend the paragraph at Page 7, lines 5-15, as follows:

A NAS system typically supports one or more file-access protocols such as NFS, WebNFS, and CIFS. "NFS" is an abbreviation for "Network File System". "CIFS" is an abbreviation for "Common Internet File System". NFS was developed by Sun® Microsystems, Inc. "WebNFS" is designed to extend NFS for use in the Internet, and was also developed by Sun Microsystems. CIFS is published as X/Open CAE Specification C209, copies of which are available from X/Open. These protocols are designed to enable a client to access remotely-stored files (or, equivalently, stored objects) as if the files were stored locally. When these protocols are used in a NAS system, the NAS controller function 230 is responsible for mapping requests which use the protocols into requests to actual storage, as discussed above. (Details of these file access protocols are not deemed necessary to an understanding of the present invention, and will not be described in detail herein.)

Please amend the paragraph beginning on Page 12, line 18 and ending on Page 13, line 14, as follows:

Oftentimes, content creators or those who deploy content within an enterprise have at least a general understanding or expectation of the relative popularity of the content which they create or deploy. For example, a developer who modifies the Web page accessed using

the URL (Uniform Resource Locator) ~~"www.ibm.com"~~ "ibm.com" which is the home page address for the IBM Corporation, knows that this page will be accessed quite frequently--and much more often than a page such as ~~"www.ibm.com/legal/copytrade.shtml"~~ "ibm.com/legal/copytrade.shtml", for example, which is a page providing information on copyrights and trademarks of IBM. This knowledge can be leveraged according to the present invention such that the content can be strategically placed at an appropriate location on the NAS resources. (It may also be appropriate in some cases to store content at more than one location on the NAS resources.) Popularity information provided by a content developer or deployer may be useful for initially placing the content. In addition, information which is subsequently learned by the developer or deployer can be used to request changes to content placement. (Note that while information for initial content placement and/or for changing content locations can be obtained from a human user based upon anticipated popularity, the human user may also specify this information based upon usage metrics which have been gathered from actual run-time data, or a combination of these types of information may be used.)

Please amend the paragraph beginning on Page 20, line 15 and ending on Page 21, line 21, as follows:

Use of the HTTP header syntax, as illustrated in FIG. 7A, enables usage metrics for any type of content object to be transmitted using a single syntax form. Assuming that an HTTP GET request such as "GET ~~http://www.ibm.com/samplePage.html HTTP/1.1~~ ibm.com/samplePage.html HTTP/1.1" is received at a Web server for which content request tracking is operating, the usage metrics of this object are updated. At some point, suppose an HTTP POST/PUT request message is sent to the Web server (which is preferably running on the NAS; see element 330 of FIG. 3) to indicate this object's popularity. The request header shown in FIG. 7A indicates the following information: (1) the content of this message pertains to relative location "samplePage/index.html", and this message is encoded using HTTP 1.1 (see element 710); (2) the content type, in this example, is "text/html" (see element 712); and (3) the usage metric for the corresponding object is, for this example, 173 (see element 714). The units of measure or other interpretation of the usage metric value may vary from one implementation to another without deviating from the scope of the present invention. Assume for purposes of illustration that the number 173 in the examples of FIGS.

7A through 7E represents 173 references to the corresponding object within the measurement interval of interest. The "UsageMetric" header shown at 714 of FIG. 7A is an example of the header syntax that the content servers which supply usage metrics (such as Web servers and/or content management systems) may generate, and that the code residing at the special control URL may search for in the metric information created by those servers, according to preferred embodiments of the present invention. Alternatively, other names for this header might be used, or individual headers might be used to separately convey factors which together comprise the overall usage metrics (such as a header for the number of access requests, a header for the length of the measurement interval or perhaps separate headers for the start and end of the measurement interval, and so forth). In this latter case, the NAS system may store these values separately upon receipt (e.g. for later use by an algorithm adapted to using those values as input parameters), or might use them to compute a result to be used in deployment determinations and then store the computed result.

Please amend the paragraph beginning on Page 22, line 17 and ending on Page 23, line 14, as follows:

Markups in other markup language objects, such as HTML and XML, may be used as alternatives to the format shown in FIG. 7A. For HTML, one example of an alternative format uses the "HTTP-EQUIV" attribute on a "META" tag, as shown at 720 in FIG. 7B. In this example, the syntax "UsageMetric" has been used as the value of the HTTP-EQUIV attribute to name the usage metric, thereby conveying information which is analogous to element 714 in FIG. 7A. A META element (such as element 720 in FIG. 7B) may be used to identify properties of a document. An HTTP-EQUIV attribute on a META tag may be used in markup language documents to explicitly specify equivalent information that an HTTP server should convey in the HTTP response message with which the document is transmitted. Information on the META tag can be found in Request For Comments ("RFC") 2518 from the Internet Engineering Task Force, which is entitled "HTTP Extensions for Distributed Authoring--WEBDAV" (February 1999), as well as on the Internet at location ~~http://wdvl.comlAuthoring/H-TML/Head/Meta/HTTP.html~~ wdvl.comlAuthoring/H-TML/Head/Meta/HTTP.html. Location ~~www.webdav.org~~ webdav.org on the Internet also provides general information about the initiative for extending the HTTP protocol to include distributed authoring and versioning syntax. Use of WebDAV requests enables accessing the

meta-data property information remotely, such that the NAS system may query the usage metric information stored at/by the Web server and/or CMS, as an alternative to using the server push techniques illustrated by FIGS. 8 and 9 (described below). (Refer to the discussion of FIGS. 7F and 7G, below, for examples of using WebDAV for obtaining usage metric information.)

Please amend the paragraph beginning on Page 23, lines 15-20, as follows:

Another example of an alternative format for use with HTML documents is the META tag with a "NAME" attribute, rather than an HTTP-EQUIV attribute. This alternative is illustrated in FIG. 7C at element 730. The NAME attribute on a META tag or element identifies a property name, and the VALUE attribute then specifies a value for that named property. For more information on use of the NAME attribute, refer to RFC 2518 or to ~~http://wdvl.com/Authoring/HTML/Head/Meta~~ wdvl.com/Authoring/HTML/Head/Meta on the Internet.

Please amend the paragraph beginning on Page 24, line 18 and ending on Page 25, line 16, as follows:

Referring now to FIGS. 7F and 7G, an example is illustrated of obtaining usage metric meta-data on request by issuing a WebDAV query and receiving a response thereto. When using WebDAV, it is likely that the hosting servers, such as Web server 140 in FIG. 1, which share a URL namespace are able to share meta-data properties about the namespace (for example, by way of a content manager). The usage metric property for a given object might then be multi-valued, wherein a given one of the values represents the usage accumulated by one of the hosting servers for that object in the URL namespace. In this case, a single WebDAV request might retrieve usage metrics for all of the servers capable of serving the content (e.g. all the servers in a server farm), as illustrated in FIG. 7F. Or, a single WebDAV request might be formulated such that it retrieves usage metrics for all of the pages served by a particular server. (This alternative has not been illustrated, but preferably comprises specifying a wildcard in the object identifying information.) The example query in FIG. 7F requests stored information that uses the "UsageMetric" property name (see element 765) associated with the content ~~"www.ibm.com/legal/copytrade.shtml"~~ "ibm.com/legal/copytrade.shtml" (see element 760), which is element 620 of FIG. 6. The

response message illustrated in FIG. 7G shows that values for this property are being provided from three Web servers "wa.foo.bar", "wb.foo.bar", and "wc.foo.bar" (see element 770), where these three Web servers may correspond to the servers 140 shown in FIG. 1. As shown in the example, the first of these servers has accumulated a total of 12 requests for this page; the second has accumulated a total of 4 requests; and the third has accumulated a total of 5 requests.